arXiv:1512.01979v1 [math.NA] 7 Dec 2015

Hyperspectral Chemical Plume Detection Algorithms Based On Multidimensional Iterative Filtering Decomposition

Antonio Cicone* Jingfang Liu Haomin Zhou†

December 8, 2015

Abstract

Chemicals released in the air can be extremely dangerous for human beings and the environment. Hyperspectral images can be used to identify chemical plumes, however the task can be extremely challenging. Assuming we know a priori that some chemical plume, with a known frequency spectrum, has been photographed using a hyperspectral sensor, we can use standard techniques like the so called matched filter or adaptive cosine estimator, plus a properly chosen threshold value, to identify the position of the chemical plume. However, due to noise and sensors fault, the accurate identification of chemical pixels is not easy even in this apparently simple situation. In this paper we present a post–processing tool that, in a completely adaptive and data driven fashion, allows to improve the performance of any classification methods in identifying the boundaries of a plume. This is done using the Multidimensional Iterative Filtering (MIF) algorithm [2, 3], which is a non–stationary signal decomposition method like the pioneering Empirical Mode Decomposition (EMD) method [9]. Moreover, based on the MIF technique, we propose also a pre–processing method that allows to decorrelate and mean–center a hyperspectral dataset. The Cosine Similarity measure, which often fails in practice, appears to become a successful and outperforming classifier when equipped with such pre–processing method. We show some examples of the proposed methods when applied to real life problems.

1 Introduction

Chemical plumes resulting from natural or anthropogenic emissions in the atmosphere can be unexpected and toxic. The detection and classification of such plumes in an efficient way would reduce the risk of harmful exposures [5, 13].

*Marie Curie fellow of the INdAM, DISIM, Universitá degli Studi dell'Aquila, via Vetoio 1, 67100, L'Aquila, Italy.

†School of Mathematics, Georgia Institute of Technology, 686 Cherry St., Atlanta, GA 30332, USA.

Recent advances in the technology provided hardware capability to easily measure the reflection energy from a large region and to generate hyperspectral images that can be used to spot and classify chemical plumes. A hyperspectral image is, in fact, a picture taken over a large number of frequencies by means of special sensors. The data produced is a hypercube whose layers are 2D images each corresponding to a different frequency. Hence each pixel of the hypercube contains a spectral signature of the corresponding physical region [1]. Given a hyperspectral image of an area where a chemical plume might be present, the plume detection problem can be divided into three subproblems. Namely: observing the presence of a plume, recognizing the chemical or chemicals contained in it, and classifying the pixels of the hyperspectral image.

Solving entirely and accurately the whole plume detection problem is not an easy task. For this reason in the end of 2012 Dimitris Manolaski and its research group at MIT Lincoln Laboratory launched the Chemical–Detection Algorithm Challenge with the idea of spurring the scientific community to solve a single subproblem. In particular they were interested in the classification stage. They gave four test datasets, each of them containing a hypercube of an area where some known chemical was released in the atmosphere, the signature of the chemical released, and the ground truth classification for each pixel. Plus they provided two blind datasets where the plume position was unknown, but it was known that a specific chemical, with a given signature, was present. The goal of the challenge was to devise a classification technique able to outperform the ones available in the literature.

Motivated by this challenge we developed, first of all, a post–processing technique, based on the Multidimensional Iterative Filtering (MIF) method [2, 3], which improves the classification accuracy of any given classifier. The MIF algorithm allows to decompose a nonstationary signal, associated with a linear or nonlinear system, into simpler components on which a time–frequency analysis can be performed. The development of methods for the decomposition of non–linear signals is a new and fast growing research area in Signal Processing. The interest in these kind of algorithms has been sparkled by the publication in 1998 of the breakthrough paper by Huang et al. [9], in which the first method of its kind, the so called Empirical Mode Decomposition (EMD) method, was proposed. The advantage of using methods like EMD or MIF versus other classical techniques available in the literature for the decomposition and time–frequency analysis of a signal, like for instance the Windowed Fourier Transform or the wavelet transform, is that these newly developed methods are designed to handle naturally non–stationary signals originated by nonlinear systems like most of the real life signals. Furthermore these techniques are completely data driven, so there is no need to make any a priori assumption on the data we want to decompose.

We note that there are many other EMD–like algorithms, such as Synchrosqueezed Wavelet Transform [4], the sparse time–frequency representation [7, 8], the Partial Differential Equation Transform [16, 17], and the ensemble EMD (EEMD) [18, 19]. They all aim at handling non-stationary signals associated with nonlinear systems. We select MIF because it is easy to implement,

ready to be applied to 2D signals, computationally fast and most importantly it has been proved to be convergent under mild conditions, at least when applied to 1D signals.

The other main result of this paper is the development of a pre–processing algorithm, also based on MIF, which allows to decorrelate and mean–center a raw hypercube allowing to use the naive Cosine Similarity (COS) measure, also known as the Spectral Angle Mapper, as a classification technique. In fact, thanks to such a pre–processing technique the COS measure produces a classification that is always comparable and sometimes outperform the ones produced by other commonly used classifiers. The advantage of coupling the COS classifier with the proposed pre–processing method is that in this way there is no more need to estimate, unlike any other commonly used classification method, the average signature and the covariance matrix of pixels not belonging to the plume.

The rest of this paper is organized as follows: in Section 2 we formalize the problem and briefly review some of the most commonly used classification methods. In Section 3 we introduce the MIF algorithm as a decomposition algorithm for one and higher dimensional signals. In Section 4 we present the proposed pre– and post–processing methods based on the MIF algorithm. Their performance are shown by means of examples in Section 5.

2 The plume detection problem and classification algorithms

A hyperspectral image is a three dimensional array, whose entries are real non-negative values, produced using hyperspectral sensors. While a standard camera sensor only captures three frequencies, the so called RGB channels, a hyperspectral sensor captures several frequencies channels at once [14]. The output of such sensors can be represented as $F \in \mathbb{R}^{h \times v \times d}$ where d is the number of frequency channels and $h \times v$ is the number of pixels in the image: to each pixel p_{ij} it corresponds a d-dimensional vector $s_{ij} \in \mathbb{R}^d$ representing its spectral signature. Assuming we are given a hyperspectral image of an area where some chemical, potentially toxic, has been released and whose signature $s_c \in \mathbb{R}^d$ is given, the problem under study is the classification of pixels in F as belonging or not to the chemical plume. In particular we want to assign to each pixel a score value in [0, 1] where 0 corresponds to a chemical free pixel and 1 to a pixel containing for sure the chemical we are looking for.

Among the classification algorithms that can be used to tackle this problem the most commonly used are the Matched Filter (MF), the Normalized Matched Filter (NMF) which is also known as Adaptive Cosine or Coherence Estimator (ACE), and the plain Cosine Similarity (COS) measure also known as the Spectral Angle Mapper [12, 11].

Assuming we have a single chemical released in the atmosphere whose spectral signature is s_c, given a pixel p with signature s its classification score value

$y(s)$ can be computed in one of the following ways:

$$y(s)_{COS} = \frac{\left(s^T s_c\right)^2}{s^T s \; s_c^T s_c} \tag{1}$$

$$y(s)_{MF} = \frac{\left[(s-\mu_b)^T \Sigma_b^{-1} s_c\right]^2}{s_c^T \Sigma_b^{-1} s_c} \tag{2}$$

$$y(s)_{ACE} = \frac{[(s-\mu_b)^T \Sigma_b^{-1} s_c]^2}{s_c^T \Sigma_b^{-1} s_c (s-\mu_b)^T \Sigma_b^{-1} (s-\mu_b)} \tag{3}$$

Where $\mu_b \in \mathbb{R}^d$ is the mean signature over all the pixels in F which are outside the plume, called for simplicity background pixels, and $\Sigma_b \in \mathbb{R}^{d \times d}$ is the covariance matrix computed considering each pixel in the background as an observation and each frequency as a variable.

We observe that if we consider the whitening and mean–centering transformation

$$\widetilde{s} = \Sigma_b^{-1/2}(s-\mu_b) \quad \text{and} \quad \widetilde{s}_c = \Sigma_b^{-1/2} s_c \tag{4}$$

then the ACE and COS classifiers produce the very same score values

$$y(\widetilde{s})_{ACE} = \frac{\left(\widetilde{s}^T \widetilde{s}_c\right)^2}{\widetilde{s}^T \widetilde{s} \; \widetilde{s}_c^T \widetilde{s}_c} = y(\widetilde{s})_{COS} \tag{5}$$

Our goal in using all these classifiers is to properly distinguish between pixels containing the known chemical and the background pixels.

Given the scores $y(s)$ of the signatures of the pixels of a hyperspectral image F, for each value of a threshold $\tau \in [0, 1]$ we can classify the pixels into four groups. If the pixel does contain the chemical and it receives a score higher than the threshold τ, it is counted as a true positive, while if the assigned score is lower than τ it is counted as a false negative. The other cases are, if the pixel does not contain the chemical and it receives a score lower than τ, it is counted as a true negative otherwise is a false positive.

To quantify the performance of a classifier the receiver operating characteristic (ROC) curves are commonly used [6].

A ROC curve is plotted in the Cartesian plane where on the vertical axis we have the ratio of true positive out of the total number of pixels that do contain the chemical, called True Positive rate (6), and the horizontal axis represents the ratio of false positive out of the total number of pixels that do not contain the given chemical, called False Positive rate (7). Both true positive and false positive rate are computed for different values of the threshold τ.

$$\text{True Positive rate (TPr)} = \frac{\text{True Positive}}{\text{True Chemical Pixels}} \tag{6}$$

$$\text{False Positive rate (FPr)} = \frac{\text{False Positive}}{\text{True Background Pixels}} \tag{7}$$

The ROC curve is a non–decreasing curve that goes from $(0,0)$ to $(1,1)$ as we vary the threshold τ. The classification produced using a random guess has a corresponding ROC curve which is the straight line connecting $(0,0)$ and $(1,1)$. The larger the area under the ROC curve (AUC), the better the performance of the classifier. For more details on ROC curves we refer the reader to [6].

In the next Section we introduce the Multidimensional Iterative Filtering method which we use to devise the proposed pre– and post–processing techniques.

3 Multidimensional Iterative Filtering (MIF)

Finding features, structures or quasiperiodicities in a non–stationary signal generated by a non–linear system can be quite challenging. However many real life problems require to handle and analyze such signals. We can think, for instance, to finance or climatic studies where the identification of recurrent patterns in data collected over time would be extremely valuable. For this reason in the last two decades many tools have been devised to deal with such signals. The goal is to decompose a signal into a finite number of simpler components called intrinsic mode functions (IMFs) plus a trend. Intuitively speaking, an IMF is a function oscillating around zero. Huang et al. in [9] defined an IMF as a function such that: in its entire domain the number of extrema and the number of zero crossings must either equal or differ at most by one; at any point, the mean value of the upper envelope defined by connecting its local maxima with splines and the lower envelope obtained by connecting its local minima is zero. From these intuitive definition it follows that sinusoidal functions are proper IMFs, but IMFs are not limited to those functions. We can have changes in the amplitude and/or the frequency as well as discontinuities in an IMF or in its derivatives.

There are many algorithm developed to decompose a signal into IMFs. The commonly used ones work either by iteration or by optimization. The pioneer work, called the Empirical Mode Decomposition (EMD) method published by Huang et al. in 1998 [9], has an iterative structure, the same as the Iterative Filtering method developed by Zhou et al. [10]. While examples of methods which are based on optimization are the Sparse Time–Frequency Representation by Hou et al. in [7, 8] as well as the Synchrosqueezed Wavelet Transforms proposed by Daubechies et al. in [4]. For an overview on such techniques we refer the reader to [2].

In this paper we focus on the Multidimensional Iterative Filtering algorithm [10, 3] (MIF), which is an extension to higher dimensions of the Iterative Filtering method that is known to be convergent under mild assumptions, at least for 1D signals, and numerical experiments showed its stability [2]. We point out here that the other known iterative method, the EMD algorithm suffers of known instabilities and its convergence is still an open problem. While a variant of the EMD method, called Ensemble EMD (EEMD), allows to solve the instability issues and it has been extended to higher dimensions [19], its convergence

remains an open problem since EEMD is based on the repeated usage of the EMD algorithm.

In the MIF algorithm[1] each IMF is produced convolving iteratively the signal with a low pass filter $w(\mathbf{t})$, $\mathbf{t} \in \Omega \subset \mathbb{R}^k$ like, for example, a Fokker–Planck filter [2] which has the nice property of being compactly supported and smooth on its entire domain. In particular, given the signal $s(\mathbf{x}), \mathbf{x} \in \mathbb{R}^k$, the MIF algorithm works as follows

MIF Algorithm IMF = MIF(s)

IMFs = {}
while the average number of extrema of $s \geq 2$ **do**
  compute the filter support Ω for s
  $s_1 = s$
  while the stopping criterion is not satisfied **do**
    $s_{n+1}(\mathbf{x}) = s_n(\mathbf{x}) - \int_\Omega s_n(\mathbf{x}+\mathbf{t}) w(\mathbf{t}) \mathrm{d}^k\mathbf{t}$
    $n = n + 1$
  end while
  IMFs = IMFs $\cup \{s_n\}$
  $s = s - s_n$
end while
IMFs = IMFs $\cup \{s\}$

Regarding the selection of the filter support Ω we want to base it adaptively on the signal. Following what proposed in [10] for the 1D case where the length of the support is computed as $4N/K$, where N is the number of sample points in the 1D signal and K is the number of its extrema, we compute for each dimension of the signal the average support length. Then we use this information to either find the radius of a spherical support or to identify the radii of an ellipsoidal $\Omega \subset \mathbb{R}^k$.

For the stopping criterion there are several possibilities. One way of doing it, as explained for the 1D case in [2, 9], is to consider the relative change

$$SD = \frac{\|s_{n+1}(\mathbf{x}) - s_n(\mathbf{x})\|_{l^2}}{\|s_n(\mathbf{x})\|_{l^2}}. \tag{8}$$

and discontinue the inner loop as soon as this relative change is less than a given threshold. The value of SD is usually set around 0.001 in our experiments.

The algorithm stops as soon as the residual $r(t)$ has a dimension which contains at most one extrema. So in the end the given signal is decomposed as $s(\mathbf{x}) = \sum_{j=1}^{m} \mathrm{IMF}_j(\mathbf{x}) + r(\mathbf{x})$, where m is the number of IMFs produced by the algorithm.

[1] The MIF code can be found at `www.cicone.com`

4 Pre– and post–processing algorithms based on MIF

In this section, we propose both a post– and a pre–processing technique based on the Multidimensional Iterative Filtering (MIF) algorithm.

The post–processing method aims to increase the performance of any given classifier. Assuming we are dealing with a chemical plume which is not punctiform, so that at least a few pixels can be classified as containing a certain chemical, the performance of a classifier are effected by noise or single sensor misfunctioning [11]. The idea, as pointed out by Wager and Walther in [15], is that pixels classified as containing the chemical we want to detect have, by assumption, to be surrounded by other pixels containing that very same chemical. We can improve the classification through the reinforcement of the spatial correlation. To do so, using the MIF method, we decompose the non–stationary and bidimensional signal, produced by a classifier, into high and low frequency components in an adaptive and data driven fashion. Then, we remove the high frequencies components, which do not contain information related with the plume, preserving instead the low frequency components. In doing so we reduce the rate of misclassified pixels improving the performance of any possible classification algorithm, not only the ones reviewed in this work. The choice of the MIF method is very natural since, unlike any other technique available in the literature, this method is designed to handle naturally non–stationary signals, it does not require to make any a priori assumption on the signal itself and its convergence has been proved when applied to 1D signals.

Post–processing method (PostP) Given the classification matrix $C \in \mathbb{R}^{h \times v}$, produced by a given classifier, we decompose C using the MIF method into its first IMF I_1 and a remainder/trend R. We then remove the first IMF from C producing the new classification $\widehat{C} = C - I_1 \equiv R$. We represent this method using the operator $\text{PostP} : \mathbb{R}^{h \times v} \mapsto \mathbb{R}^{h \times v}$, $C \to \widehat{C} = \text{PostP}(C)$.

We observe that, since the convergence of this method applied to the decomposition of 2D signals is still an open problem, we checked the meaningfulness of the outcome of such a post–processing using the 1D convergent MIF method . We consider, in fact, the classification scores first along each row and then along each column of the classification matrix, and vice versa. The results of these two procedures prove to be close to the one obtained post–processing directly the 2D classification using the 2D MIF method.

The other method we present here is a pre–processing technique. The goal is to devise a completely data driven decorrelation and mean–centering technique that can be applied to a hyperspectral dataset before classifying it with the plain Cosine Similarity (COS) measure. The idea behind this pre–processing method comes from the observations, we made previously, regarding the ACE and COS classifiers. In particular we know that ACE requires the a priori knowledge of the mean signature and covariance matrix of the background and we observed that if we can somehow whiten and mean–center the hyperspectral

image pixel signatures then the COS and ACE classifiers produce the very same classification scores (5). To achive this goal we propose the following procedure. First of all we subtract from each pixel signature s the average signature μ of the entire hyperspectral dataset. This is done in order to remove possible artifacts introduced uniformly in all the pixel signatures by the sensor or the device used to capture the hyperspectral image. Then we decompose each pixel signature, using the MIF method, into IMFs and a trend. Hence we subtract such trend from the original signature. The result is a hyperspectral dataset which is mean–centered and decorrelated in a local way and it is ready to be classified using the plain Cosine Similarity (COS) measure.

Pre–processing method (PreP) Given a hypercube F and the signature $s_{ij} \in \mathbb{R}^d$ of the (i, j) pixel, we first compute $\mu \in \mathbb{R}^d$, mean signature of all the pixels in F, and we subtract it from each pixel signature to produce $\widetilde{s}_{ij} = s_{ij} - \mu$. Then using the MIF method we decompose $\widetilde{s}_{ij}$ into its IMFs $\{I_n\}_{n=1,\dots N}$, for some $N \in \mathbb{N}$, and a trend r so that $\widetilde{s}_{ij} = \sum_{n=1}^{N} I_n^{(i,j)} + r$. Finally we remove from the signature $\widetilde{s}_{ij}$ the remainder/trend r. So each pixel signature becomes $\widehat{s}_{ij} = \sum_{n=1}^{N} I_n^{(i,j)}$. We can summarize the entire procedure using the operator PreP : $\mathbb{R}^{h \times v \times d} \mapsto \mathbb{R}^{h \times v \times d}$, $F \to \widehat{F} = \mathrm{PreP}(F) = [\widehat{s}_{ij}]_{i,\,j}$.

The proposed pre–processing technique can be applied to hyperspectral images. It allows, in particular, to decorrelate and mean–center the hypercubes. Since we simply decorrelate and we do not whiten the dataset we do not expect COS and ACE classifications to match in general. However from the experiments we run we observe that COS equipped with such a pre–processing method proves to have performance that are always similar and sometimes even better than the one of the ACE classifier. Therefore, from a COS classification point of view, the global and local trend removal performed in the proposed pre–processing method proves to be an equivalent procedure to the mean–centering and whitening of the data, as shown in the next section.

5 Examples

In this section we show the performance of the proposed pre– and post–processing methods when applied to the plume detection problem. We consider the datasets provided for the 2012 DTRA Chemical Detection Challenge[2]. They contain both real world and synthetic hyperspectral images, and the signatures of the chemical contained in the plume are always included. In this work we focus on the real world datasets, since the analysis of the synthetic ones appear to be trivial in general. We test our methods first on a dataset where the ground truth is known, since the actual position of the chemical plume is given, and

[2]The datasets can be obtained from the National Science Foundation (NSF). However, they are provided only to researchers and teams supported by NSF itself.

then on a blind one that have been provided for the contest. With the former we can evaluate the performance of each classifier using ROC curves.

We point out here that the ground truths provided in the datasets for the 2012 DTRA Chemical Detection Challenge assume a classification of the hyperspectral image pixels into three groups: "inside the plume", "outside the plume" and "close to the boundaries of the plume", as shown, for instance, in Figure 1b. As explained in Section 2, in this field of research the goal is to devise a method able to classify pixels using just two classes: pixels containing the known chemical and background pixels. The reason why the ground truth includes the additional class of pixels close to the boundaries of the plume is because classifying such pixels as either inside or outside the plume is a hard problem. Therefore, following what suggested by Manolakis and his research group for the 2012 DTRA Chemical Detection Challenge, we use the ground truth to check the performance of a classifier only on pixels classified as either "inside the plume" or "outside the plume". This means that the ROC curves presented in this Section are built using solely the pixels that are not contained in the set "close to the boundaries of the plume". It is clear that if we use only ROC curves built in the aforementioned way we can compare the performance of different classifiers on pixels outside and inside the plume, but we are unable to compare them on pixels close to the boundaries of the plume. For this reason we show both ROC curves and the pixel classification values produced by each classifier in order to give readers a more comprehensive picture of the performance of each classifier.

In the blind case, instead, it is impossible to plot ROC curves since the ground truth is unknown. In this second case we show the performance of a classifier simply plotting the pixel classification values, also known as detection maps.

We plot the pixel classification values using "flipud" colormap, option "hot", in Matlab. In all the examples presented in the following the colormap is calibrated identically.

We would like to note that we design MIF and the pre– and post–processing techniques for hyperspectral data classifications not to compete with the existing classification methods, but to use them as complementary strategies to further improve their performances. In this paper, we present the performances of the proposed pre– and post–processing techniques when applied to standard classifiers like ACE, MF and COS. We point out that they can be easily adopted to other classifiers as well.

5.1 Case of known plume position: hypercube "Location 1 released r134a"

The first dataset we test contains the hypercube plotted in Figure 1a taken in an area where the chemical r134 has been released. The signature of this chemical is provided in the dataset and shown in red in Figure 7b. The given ground truth is shown in Figure 1b where pixels inside the plume are colored black, while ones close to the boundaries of the plume are colored orange. We observe from

this figure that the boundary region, that we are excluding in the evaluation of the ROC curves, as explained previously, is pretty wide. It becomes evident that using only ROC curves and area under the curve (AUC) values provides only partial information on the true performance of each classifier.

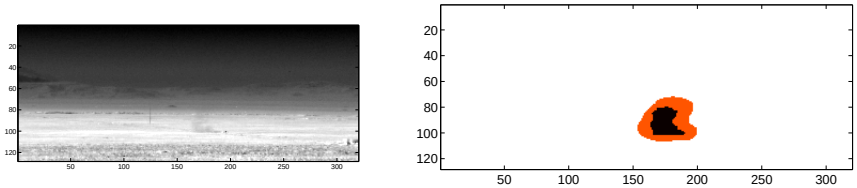

(a) Contrast–enhanced spectral–mean image using the raw data from the dataset *Location 1 released r134a*

(b) Ground truth classification. In black plume pixels, in orange pixels close to the boundaries of the plume.

Figure 1

We start applying the ACE classifier to this hypercube. The outcomes are plotted in Figure 2a where we clearly see the contribution of the noise. The corresponding ROC curve is shown in Figure 4a, its AUC value is approximately 0.98919. If we apply the post–processing method **PostP** described in the previous section we clean the classification preserving the shape of the plume as shown both in Figure 2b and by the ROC curve plotted in Figure 4a. Its AUC value is approximately 0.99508. We compare this result with the post–processing methods based on wavelet transform and on the Multi-dimensional empirical EMD (MEEMD) method, which is a multidimensional version of the ensemble EMD proposed by Huang et al. in [19]. For the first one we use the Matlab noise removal tool called *Stationary Wavelet Transform Denoising 2–D*, where we choose the Daubechies db2 wavelet to get the cleaned classification shown in Figure 3a, whose ROC curve is plotted in Figure 4a. The AUC value in this case is approximately 0.99507. While for the MEEMD method we use the code provided in [19] where we set to 10 the number of perturbations. The post–processed classification is obtained subtracting the first 2 IMFs from the original classification. The corresponding cleaned classification is depicted in Figure 3b, its ROC curve is shown in Figure 4a and the AUC value is exactly equal, up to machine precision, to the one obtained using the **PostP** approach.

We point out that, since in all the examples under study the roc curves tend to be close each other and the differences are concentrated in the top left corner of the ROC curve plots, from now on we opt to plot ROC curves using a log scale along the horizontal axis, ref. Figure 4b.

Another observation regards the shape of the boundaries produced using these post– processing techniques. As we already mentioned, the pixels close to the boundaries of the plume, shown in orange in Figure 1b, are completely excluded in the computation of the ROC curves. This is the reason why all the classifications end up having almost identical ROC curves and area under the curve values, even when they have pretty different pixel classification values as

for the case of the wavelet transform post–processing versus the one obtained with the **PostP** and the MEEMD techniques.

To have a better understanding of the true performances of these different post–processing methods we can compare directly the classifications depicted in Figures 2b, 3a and 3b with the ground truth shown in Figure 1b. From this comparison it is evident that when using the wavelet transform, while the noise is removed in a reasonable way, the shape of the plume is partially lost due to excessive blunting. Since in this context it is very important to identify the boundaries of the plume with high accuracy, the wavelet transform proves to be not the best option as a post–processing technique. More in general, in all the examples we tested, standard techniques for noise removal, like Fourier or wavelets transform, tend to have problems producing accurate boundaries. This is due to two reasons: this kind of signals are in general non-stationary, and the standard transformations often use bases that are determined a priori. Both limitations are overcome by the usage of the Multidimensional Iterative Filtering (MIF) method as a post–processing technique. In fact MIF technique can handle a non-stationary signal and does not require any a priori assumption or knowledge on it, as confirmed in this example by the shape of the plume and its boundary after MIF post–processing, ref. Figure 2b.

Regarding the results obtained via MEEMD we point out that, even if the results are comparable with the one produced with the proposed method, this approach has disadvantages. First of all it is still an open problem to prove the convergence of this technique even in one dimension. In fact in this case the method reduces to the so called EEMD technique, whose convergence is not known. Whereas we recall that, even if the convergence of the MIF method is still an open problem as well, its one dimensional version, the Iterative Filtering technique, has been proven to be convergent under mild conditions [2]. Furthermore the MEEMD proves to be much slower than the proposed method. On an average personal computer it takes around 4 seconds to produce the clean classification plotted in Figure 2b, while the MEEMD runs for more than 190 seconds on the same computer to produce the result shown in Figure 3b.

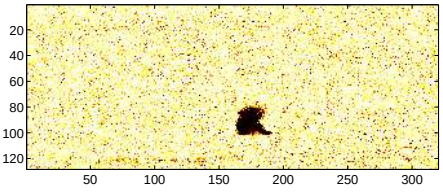

(a) Detection map of the ACE classification of the raw data

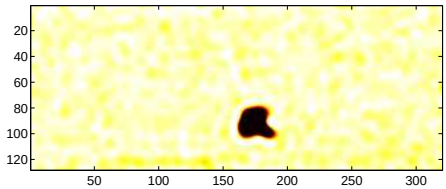

(b) Detection map after applying the **PostP** method to the ACE classification of the raw data

Figure 2

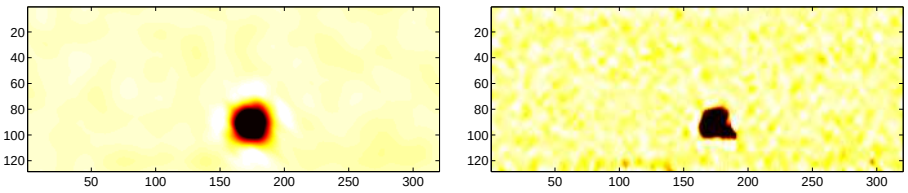

(a) Detection map of the Wavelets post–processed ACE classification of the raw data

(b) Detection map of the MEEMD post–processed ACE classification of the raw data

Figure 3

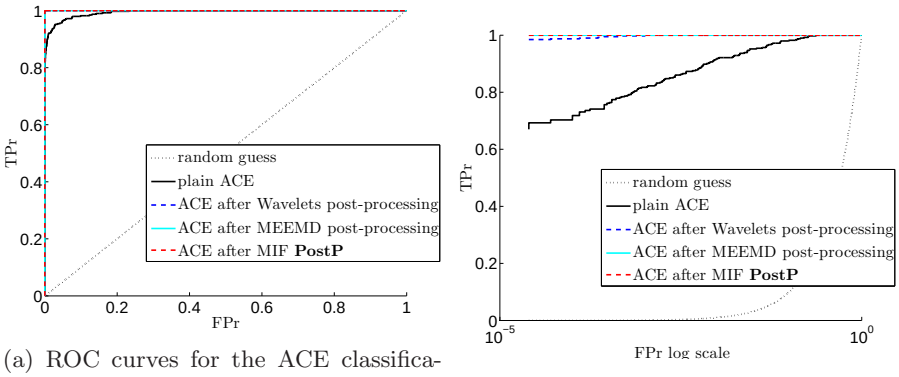

(a) ROC curves for the ACE classification of the hypercube *Location 1 released r134a* with and without post–processing

(b) ROC curves in log scale along the horizontal axis

Figure 4

We observe that with this dataset the MF classifier gives results similar to ACE, so we avoid presenting them and move directly to the COS classifier.

The pixel classification produced by COS applied to the raw hypercube is shown in Figure 5a where we reversed the pixel classification values. We need to do so in order to have meaningful results. In fact the corresponding ROC curve, plotted in dashed black in Figure 8a, would be below the random guess curve without the reversion. If we apply the **PostP** method to the reversed COS classification we can improve the classification performance. This is testified by the increasing of the area under the ROC curve, ref. ROC curve plotted in solid black in Figure 8a. However, even with post–processing, the COS classification of the raw hyperspectral image still needs to be reversed. To solve this issue we devise and apply the pre–processing method **PreP** based on MIF algorithm and described in the previous section. The COS classification of the pre–processed hypercube is shown in Figure 5b. Its ROC curve, plotted in dashed blue in Figure 8a, is now directly above the random guess curve without any need of reversing the classification. Nevertheless its performances are worse than the

ones produced by the reversed COS classification of the raw dataset, the dashed black curve is clearly above the solid blue one. However if we apply the **PostP** method to the COS classification of the pre–processed hypercube we get better performance than any other COS classification with or without post–processing. In fact the solid blue ROC curve is above all the others in Figure 8a. We can also try to apply the wavelet transform as a denoising post–processing method. If we plot the corresponding ROC curve, its performances are slightly worse than the ones obtained using the proposed method, the solid magenta curve is slightly lower than the solid blue one in Figure 8a. Furthermore the boundaries end up being excessively blunted, Figure 6b, for the reasons explained above, while the proposed post–processing technique allows to produce boundaries that look more natural, ref. Figure 6a.

So, overall, the best performances are obtained from the combination of the pre– and post–processing techniques, both based on MIF, applied to the COS classifier. We point out that the range of values produced by the COS classifier applied to the raw dataset is in the narrow interval [0.8, 0.84] after the reversion, while the one corresponding to the COS classification of the pre–processed dataset is in the wider interval [0, 0.5] without any need of reversing them. The performance of the COS classifier equipped with the pre–processing method becomes comparable with the one of ACE and MF algorithms.

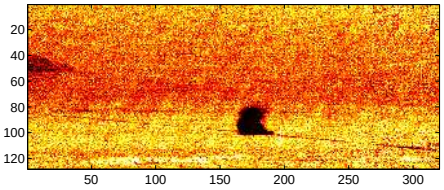

(a) Detection map of the reversed COS classification of the raw data

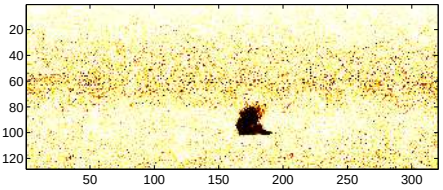

(b) Detection map of the COS classification of the pre–processed data

Figure 5

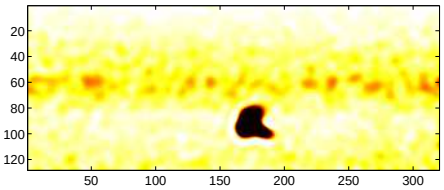

(a) Detection map of the **PostP** post–processed COS classification of the pre–processed data

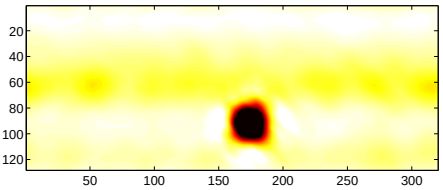

(b) Detection map of the Wavelets post–processed COS classification of the pre–processed data

Figure 6

A question that naturally arises at this point is why the pre–processing

method works so well in enhancing the COS classification. The reason is evident from Figures 7a and 7b. Removing the global and local trend from the signature of each pixel allows to properly measure the angle between the pixels and the chemical signatures. Furthermore we observe that all the datasets provided for the aforementioned challenge always have the property that the signature of the chemical seems to be subtracted from the signature of the pixels contained in the plume. It remains an open problem to understand the reason of such a behavior of the signatures. From this last observation it is evident that only after removing the trend of the pixel signatures the direct comparison of the signatures made by the COS classifier produces a meaningful classification. We observe that the combination of a pre– and post–processing of the COS classification works well in increasing the performance of this classifier for all the datasets provided for the 2012 DTRA Chemical Detection Challenge.

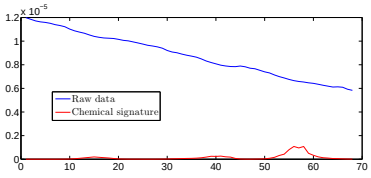

(a) Raw spectral signature of a pixel inside the plume and the chemical r134 signature

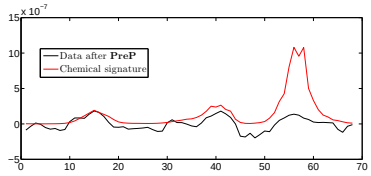

(b) Pre–processed spectral signature of a pixel inside the plume

Figure 7

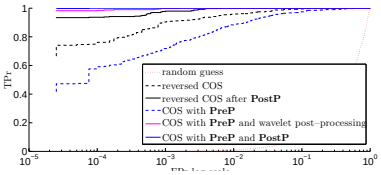

(a) ROC curves for the COS classification of the hypercube *Location 1 released r134a*

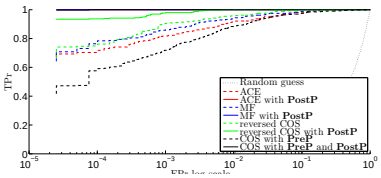

(b) ROC curves comparison for the hypercube *Location 1 released r134a*

Figure 8

5.2 Blind case: hypercube "Location 2 released sf6 blind"

As another example we consider one of the blind datasets provided for the challenge. The hypercube image is shown in Figure 9a. It is known that this image contains a plume of the sf6 chemical, whose signature is given.

If we apply the ACE classifier to the hypercube as it is, we obtain the classification depicted in Figure 9b. We can apply the **PostP** post–processing method

to reenforce the spatial correlations, ref. Figure 10a. The results produced by the Matlab noise removal tool called *Stationary Wavelet Transform Denoising 2–D*, with the choice of the Daubechies db2 wavelet, are depicted in Figure 10b. Similar results we get from the post–processing of the MF classification, ref. Figure 11a. Using the proposed **PostP** technique, Figure 11b, we clean in a more effective way than with a standard technique, ref. Figure 12a.

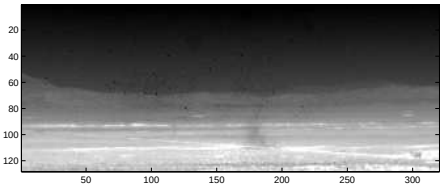

(a) Contrast–enhanced spectral–mean image using the raw data *Location 2 released sf6 blind*

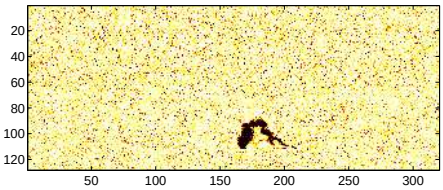

(b) Detection map of the ACE classification of the raw data

Figure 9

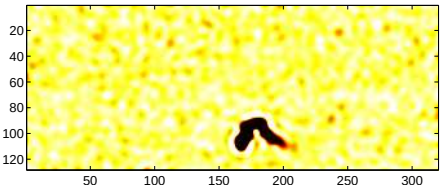

(a) Detection map of the MIF post–processed ACE classification of the raw data

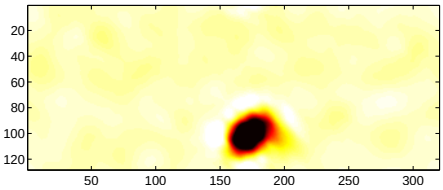

(b) Detection map of the Wavelets post–processed ACE classification of the raw data

Figure 10

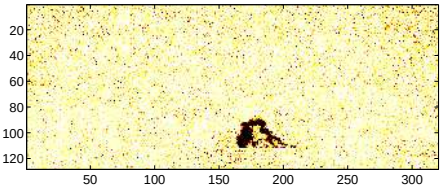

(a) Detection map of the MF classification of the raw data

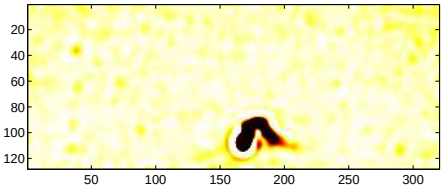

(b) Detection map of the MIF post–processed MF classification of the raw data

Figure 11

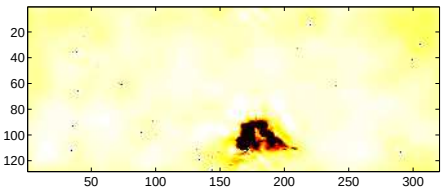

(a) Detection map of the Wavelets post–processed MF classification of the raw data

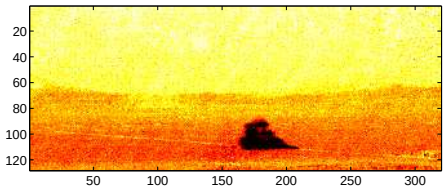

(b) Detection map of the reversed COS classification of the raw data

Figure 12

With the plain COS classifier, as it was for the previous dataset, we need to reverse the classification in order to produce meaningful results, Figure 12b. If we apply the proposed **PreP** pre–processing method to the hypercube we can remarkably increase the performance of the classifier as shown in Figure 13a so that there is no more need to reverse the classification and its performance become better than the one of the other classifier tested. In fact the COS method equipped with the proposed pre–processing method is producing a more natural plume shape than the one produced by any other classification techniques tested. So the pre–processing technique allows to transform a known failing classification method into what appears to be an extremely successful classification algorithm. We observe that the numerical range of the COS classification of to the raw hypercube spans approximately the interval 0.92 and 0.94, while the one of the pre–processed hypercube is in between 0 and 0.5.

Also in the case of the COS classifier the application of the **PostP** post–processing technique helps in cleaning and improving the classification, Figure 13b.

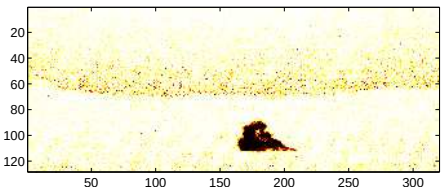

(a) Detection map of the COS classification of the pre–processed data

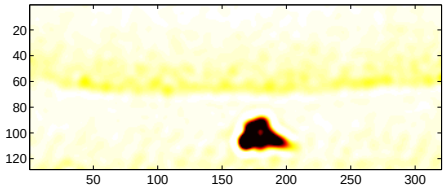

(b) Detection map of the post–processed COS classification of the pre–processed data

Figure 13

6 Conclusions

Inspired by the DTRA 2012 Chemical–Detection Algorithm Challenge posted by Dimitris Manolakis and his research group at MIT Lincoln Laboratory, in

this paper we tackle the problem of the accurate detection of the boundaries of a chemical plume hidden in a hyperspectral image by means of classification algorithms. In particular we devise a post– and a pre–processing algorithm aimed to improve the performance of known classifiers.

The post–processing technique we present in this paper is based on the Multidimensional Iterative Filtering method [3] which is an algorithm for the decomposition of non–stationary signals. This post–processing technique is a uniform procedure which allows to improve the performance of any classifier used in the chemical plume detection problem by means of an adaptive and data driven cleaning and smoothing of its classification. The choice of the MIF method appears to be ideal, since it is an algorithm developed to handle naturally non–stationary signals, which is the case for almost any real life signal. Furthermore MIF algorithm does not require to make any a priori assumption on the kind of data we are analyzing as for other techniques like wavelet transform and similar methods. We test this approach on the datasets provided for the aforementioned challenge, and we show some of these results in the Examples section.

Secondly we propose a pre–processing technique which, independently from the dataset under study, allows to decorrelate and mean–center the data. The goal of this procedure is to make the score values produced by the COS classifier equal to the ones produced by the ACE classifier, as shown in (5). This is made possible thanks to the employment of the Multidimensional Iterative Filtering method which allows to decompose a signal, in particular a non–stationary one, in an adaptive and data driven way. From the tests run, and partially shown in the examples section, we notice that the chemical plume boundaries detected using the COS classifier equipped with the **PreP** pre–processing method appear to be meaningful. These results are confirmed by the ROC curves and the shape of the detected plumes. In some cases the detected boundaries prove to be more meaningful and reasonable than the ones produced using other standard classification methods. In conclusion, given the Cosine Similarity measure, which is known to be a failing classification technique, the proposed pre–processing technique allows to transform it into a successful classifier which can even outperform in some cases other commonly used classifiers.

A possible direction for future researches is in the analysis of plumes containing mixtures of two or more known chemicals. Traditionally such mixtures have been assumed to be linear, however it is known that such assumption is far from reality. New ways for tackling this problem have to be devised. We point out also that, from an algorithmic point of view, the convergence of the Multidimensional Iterative Filtering method has been proved only for the 1D case [2]. For higher dimensions its convergence has to be studied yet.

7 Acknowledgments

This work was supported by NSF Faculty Early Career Development (CAREER) Award DMS–0645266, DMS–1042998, DMS–1419027, and ONR Award

N000141310408.

Antonio Cicone acknowledges support by National Group for Scientific Computation (GNCS – INdAM) "Progetto giovani ricercatori 2014", and by Istituto Nazionale di Alta Matematica (INdAM) "INdAM Fellowships in Mathematics and/or Applications cofunded by Marie Curie Actions".

References

[1] C.-I. Chang. *Hyperspectral imaging: techniques for spectral detection and classification*, volume 1. Springer, 2003.

[2] A. Cicone, J. Liu, and H. Zhou. Adaptive local iterative filtering for signal decomposition and instantaneous frequency analysis. submitted to *Applied and Computational Harmonic Analysis*. *arXiv preprint arXiv:1411.6051*, 2014.

[3] A. Cicone, and H. Zhou. Multidimensional Iterative Filtering method for the decomposition of high–dimensional non–stationary signals. *arXiv preprint arXiv:1507.07173*, 2015.

[4] I. Daubechies, J. Lu, and H.-T. Wu. Synchrosqueezed wavelet transforms: An empirical mode decomposition–like tool. *Applied and Computational Harmonic Analysis*, 30(2):243–261, 2011.

[5] V. Farley, V. M. Chamberland, P. Lagueux, A. Vallie'res, A. Villemaire, and J. Giroux. Chemical agent detection and identification with a hyperspectral imaging infrared sensor *Proceedings of SPIE*,6739:18–34,2007.

[6] T. Fawcett. An introduction to roc analysis. *Pattern recognition letters*, 27(8):861–874, 2006.

[7] T.Y. Hou, M.P. Yan, and Z. Wu. A variant of the emd method for multi-scale data. *Adv. in Adap. Data Anal.*, 1(4):483–516, 2009.

[8] T.Y. Hou and Z. Shi. Adaptive data analysis via sparse time-frequency representation. *Adv. in Adap. Data Anal.*, 3(1):1–28, 2011.

[9] N.E. Huang, Z. Shen, S.R. Long, M.C. Wu, H.H. Shih, Q. Zheng, N.C. Yen, C.C. Tung, H.H. Liu. The empirical mode decomposition and the hilbert spectrum for nonlinear and non-stationary time series analysis. *Proceedings of the Royal Society of London. Series A*, 454(1971):903, 1998.

[10] L. Lin, Y. Wang, and H. Zhou. Iterative filtering as an alternative algorithm for empirical mode decomposition. *Advances in Adaptive Data Analysis*, 1(4):543–560, 2009.

[11] D. G. Manolakis, S. E. Golowich, and R.S. DiPietro. Long-wave infrared hyperspectral remote sensing of chemical clouds: A focus on signal processing approaches. *Signal Processing Magazine, IEEE* 31(4):120–141, 2014.

[12] D. Manolakis and G. Shaw. Detection algorithms for hyperspectral imaging applications. *Signal Processing Magazine, IEEE*, 19(1):29–43, 2002.

[13] S. Niu, S. E. Golowich, V. K. Ingle, and D. Manolakis. Hyperspectral chemical plume quantification via background radiance estimation. In *SPIE Defense, Security, and Sensing*, 874316–874316. International Society for Optics and Photonics, 2013.

[14] G. Shaw, and D. Manolakis. Signal processing for hyperspectral image exploitation. *Signal Processing Magazine, IEEE*, 19(1):12–16, 2002.

[15] S. Wager, and G. Walther. Smooth chemical Vapor detection. *talk at DTRA conference in Boulder, CO*, 2012.

[16] Y. Wang, W. Guo-Wei, and Y. Siyang Partial differential equation transformvariational formulation and Fourier analysis. *International journal for numerical methods in biomedical engineering*, 27(12):1996–2020, 2011.

[17] Y. Wang, W. Guo-Wei, and Y. Siyang Mode decomposition evolution equations. *Journal of scientific computing*, 50(3):495–518, 2012.

[18] Z. Wu and N.E. Huang. Ensemble empirical mode decomposition: A noise-assisted data analysis method. *Advances in Adaptive Data Analysis*, 1(1):1–41, 2009.

[19] Z. Wu, N.E. Huang, and X. Chen. The multi-dimensional ensemble empirical mode decomposition method. *Advances in Adaptive Data Analysis*, 1(3):339–372, 2009.